

# Go-With-The-Winner: Client-Side Server Selection for Content Delivery

Chang Liu  
University of Massachusetts,  
Amherst  
cliu@cs.umass.edu

Ramesh K. Sitaraman  
University of Massachusetts,  
Amherst  
and Akamai Technologies Inc.  
ramesh@cs.umass.edu

Don Towsley  
University of Massachusetts,  
Amherst  
towsley@cs.umass.edu

## ABSTRACT

Content delivery networks deliver much of the world's web and video content by deploying a large distributed network of servers. We model and analyze a simple paradigm for client-side server selection that is commonly used in practice where each user independently measures the performance of a set of candidate servers and selects the one that performs the best. For web (resp., video) delivery, we propose and analyze a simple algorithm where each user randomly chooses two or more candidate servers and selects the server that provided the best hitrate (resp., bitrate). We prove that the algorithm converges quickly to an optimal state where all users receive the best hitrate (resp., bitrate), with high probability. We also show that if each user chose just one random server instead of two, some users receive a hitrate (resp., bitrate) that tends to zero. We simulate our algorithm and evaluate its performance with varying choices of parameters, system load, and content popularity.

## 1. INTRODUCTION

Modern content delivery networks (CDNs) host and deliver a large fraction of the world's web content, video content, and application services on behalf of enterprises that include most major web portals, media outlets, social networks, application providers, and news channels [17]. CDNs deploy large numbers of servers around the world that can store content and deliver that content to users who request it. When a user requests a content item, say a web page or a video, the user is directed to one of the CDN's servers that can serve the desired content to the user. The goal of a CDN is to maximize the performance as perceived by the user while efficiently managing its server resources.

A key functionality of a CDN is the *server selection* process by which client software running on the user's computer or device, such as media player or a browser, is directed to a suitable server of a CDN [7]. The desired outcome of the server selection process is that each user is directed to a server that can provide the requested content with good

performance. The metrics for performance that are optimized vary by the type of content being accessed. For instance, good performance for a user accessing a web page might mean that the web page downloads quickly. Good performance for a user watching a video might mean that the content is delivered by the server at a sufficiently high bitrate to avoid the video from freezing and rebuffering [12].

Server selection can be performed in two distinct ways that are not mutually exclusive. *Network-side server selection* algorithms monitor the real-time characteristics of the CDN and the Internet. Such algorithms are often complex and measure liveness and load of the CDN's servers, as well as latency, loss, and bandwidth of the communication paths between servers and users. Using this information, the algorithm computes a good "mapping" of users to servers, such that each user is assigned a "proximal" server capable of serving that user's content [17]. This mapping is computed periodically and is typically made available to client software running on the user's computer or device using the domain name system (DNS). Specifically, the user's browser or media player looks up the domain name of the content that it wants to download and receives as translation the ip address of the selected server.

A complementary approach to network-side server selection that is commonly used is *client-side server selection* where the client software running on the user's computer or device embodies a server selection algorithm. The client software is typically unaware of the global state of the server infrastructure, the Internet, or other users. Rather, the client software typically makes future server selection decisions based on its own historical performance measurements from past server downloads. Client-side server selection can often be implemented as a plug-in within media players, web browsers, and web download managers [2].

While client-side server selection can be used to select servers within a single CDN, it can also be used in a multi-CDN setting. Large content providers often make the same content available to the user via multiple CDNs. In this case, the client software running on the user's device tries out the different CDNs and chooses the "best" server from across multiple CDNs. For instance, Netflix uses three different CDNs and the media player incorporates a client-side server selection algorithm to choose the "best" server (and the corresponding CDN) using performance metrics such as the video bitrates achievable from the various choices [1]. Note also that in the typical multi-CDN case, both network-side and client-side server selection are used together, where the former is used to choose the candidate servers from each




Figure 1: Client-side Server Selection with the Go-With-The-Winner paradigm. User  $U$  makes request to two candidate servers  $S1$  and  $S2$ . After a trial period of observing the performance provided by the candidate, the user selects the better performing server.

CDN and the latter is used by the user to pick the “best” among all the candidates.

### 1.1 The Go-With-The-Winner paradigm

A common and intuitive paradigm that is often used for client-side server selection in practice is what we call “Go-With-The-Winner” that consist of an initial *trial period* during which each user independently “tries out” a set of *candidate servers* by requesting content or services from them (cf. Figure 1). Subsequently, each user independently *decides* on the “best” performing server using historical performance information that the user collected for the candidate servers during the trial period. It is commonly implemented in the content delivery context that incorporate choosing a web or video content server from among a cluster of such servers.

Besides content delivery, the Go-With-The-Winner paradigm is also common for other Internet services, though we do not explicitly study such services in our work. For instance, BIND, which is the most widely deployed DNS resolver (i.e., DNS client) on the Internet, tracks performance as a smoothed value of the historical round trip times (called SRTT) from past queries for a set of candidate name servers. Then BIND chooses a particular name server to query in part based on the computed SRTT values [13]. It is also notable that BIND implementations incorporate randomness in the candidate selection process.

The three key characteristics of the Go-With-The-Winner paradigm are as follows.

1. *Distributed control.* Each user makes decisions in a distributed fashion using only knowledge available to it. There is no explicit information about the global state of the servers or other users, beyond what the user can infer from it’s own historical experience.
2. *Performance feedback only.* There is no explicit feedback from a server to a user who requested service beyond what can be inferred by the performance ex-

perienced by the user.

3. *Choosing the “best” performer.* The selection criteria is based on historical performance measured by the user and consists of selecting the best server according to some performance metric (i.e., go with the winner).

Besides its inherent simplicity and naturalness, the paradigm is sometimes the only feasible and robust solution. For instance, in many settings, the client software running on the user’s device that performs server selection has no detailed knowledge of the state of the server infrastructure as it is managed and owned by other business entities. In this case, the primary feedback mechanism for the client is its own historical performance measurements.

While client-side server selection is widely implemented, its theoretical foundations are not well understood. A goal of our work is to provide such a foundation in the context of web and video content delivery. *It is not our intention to model a real-life client-side server selection process in its entirety which can involve other adhoc implementation-specific considerations. But rather we abstract an analytical model that we can explore to extract basic principles of the paradigm that are applicable in a broad context.*

### 1.2 Our contributions

We propose a simple theoretical model for the study of client-side server selection algorithms that use the Go-With-The-Winner paradigm. Using our model, we answer foundational questions such as how does randomness help in the trial period when selecting candidate servers? How many candidate servers should be selected in the trial phase? How long does it take for users to narrow down their choice and decide on a single server? Under what conditions does the selection algorithm converge to a state where all users have made the correct server choices, i.e., the selected servers provide good performance to their users? Some of our key results that help answer these questions follow.

(1) In Section 2, in the context of web content delivery, we analyze a simple algorithm called GoWithTheWinner where each user independently selects two or more random servers as candidates and decides on the server that provided the best cache hit rate,. We show that with high probability, the algorithm converges quickly to a state where no cache is overloaded and all users obtain a 100% hit rate. Further, we show that two or more random choices of candidate servers are necessary, as just one random choice will result in some users (and some servers) incurring cache hit rates that tend to zero, as the number of users and servers tend to infinity. This work represents the first demonstration of the “power of two choices” phenomena in the context of client-side server selection for content delivery, akin to similar phenomena observed in balls-into-bins games [14], load balancing, circuit-switching algorithms [5], relay allocation for services like Skype [16], and multi-path communication [11].

(2) In Section 3, in the context of video content delivery, we propose a simple algorithm called MaxBitRate where each user independently selects two or more random servers as candidates and decides on the server that provided the best bitrate for the video stream, We show that with high probability, the algorithm converges quickly to a state where no server is overloaded and all users obtain the required bitrate for their video to play without freezes. Further, we

show that two or more random choices of candidate servers are necessary, as just one random choice will result in some users receiving bitrates that tend to zero, as the number of users and servers tends to infinity.

(3) In Section 4, we go beyond our theoretical model and simulate algorithm *GoWithTheWinner* in more complex settings. We establish an inverse relationship between the length of the history used for hitrate computation (denoted by  $\tau$ ) and the failure rate defined as the probability that the system converges to a non-optimal state. We show that as  $\tau$  increases the convergence time increases, but the failure rate decreases. We also empirically evaluate the impact of the number of choices of candidate servers. We show that two or more random choices are required for all users to receive a 100% hitrate. Though even if only 70% of the users make two choices, it is sufficient for 95% of the users to receive a 100% hitrate. Finally, we show that the convergence time increases with system load. But, convergence time decreases when the exponent of power law distribution that describes content popularity increases.

## 2. HIT RATE MAXIMIZATION FOR WEB CONTENT

The key measure of web performance is *download time* which is the time taken for a user to download a web object, such as a html page or an embedded image. CDNs enhance web performance by deploying a large number servers in access networks that are “close” to the users. Each server has a cache that is capable of storing web objects. When a user requests an object, such as a web page, the user is directed to a server that can serve the object (cf. Figure 1). If the server already has the object in its cache, i.e., the user’s request is a *cache hit*, the object is served from the cache to the user. In this case, the user experiences good performance, since the CDN’s servers are proximal to the user and the object is downloaded quickly. However, if the requested object is not in the server’s cache, i.e., the user’s request is a *cache miss*, then the server first fetches it from the origin, places it in its cache, and then serves the object to the user. In the case of a cache miss, the performance experienced by the user is often poor since the origin server is typically far away from the server and the user. In fact, if there is a cache miss, the user would have been better off not using the CDN at all, since downloading the content directly from the content provider’s origin would likely have been faster! Since the size of a server’s cache is bounded, cache misses are inevitable. A key goal of server selection for web content delivery is to jointly orchestrate server assignment and content placement in caches such that the cache hit rate is maximized. While server selection in CDNs is a complex process [17], we analytically model the key elements that relate to content placement and cache hit rates, leaving other factors that impact performance such as server-to-user latency for future work.

### 2.1 Problem Formulation

Let  $U$  be a set of  $n_u$  users who each request an object picked independently from a set  $C$  of size  $n_c$  using a power law distribution where the  $k^{th}$  most popular object in  $C$  is picked with a probability

$$p_k \triangleq \frac{1}{k^\alpha \cdot \mathcal{H}(n_c, \alpha)}, \quad (1)$$

where  $\alpha \geq 0$  is the exponent of the distribution and  $\mathcal{H}(n_c, \alpha)$  is the generalized harmonic number that is the normalizing constant, i.e.,  $\mathcal{H}(n_c, \alpha) = \sum_{k=1}^{n_c} 1/k^\alpha$ . Note that power law distributions (aka Zipf distributions) are commonly used to model the popularity of online content such as web pages, and videos. This family of distributions is parametrized by a Zipf rank exponent  $\alpha$  with  $\alpha = 0$  representing the extreme case of an uniform distribution and larger values of  $\alpha$  representing a greater skew in the popularity. It has been estimated that the popularity of web content can be modeled by a power law distribution with an  $\alpha$  in the range from 0.65 to 0.85 [4, 10, 9]. The user then sticks with that content and makes a sequence of requests to the set of available servers. Relating to the reality, users tend to stay with one website for a while, say reading the news or looking at a friend’s posts. Here the *whole website* is what we considered a content. We model the sequence of requests generated by each user as a Poisson process with a homogeneous arrival rate  $\lambda$ . Note that each request from user  $u$  can be sent to one or more servers selected from  $S_u \subseteq S$ , where  $S_u$  is the server set chosen by user  $u$ .

Let  $S$  be the set of  $n_s$  servers that are capable of serving content to the users. Each server can cache at most  $\kappa$  objects and a cache replacement policy such as LRU is used to evict objects when the cache is full. Given that the download time of a web object is significantly different when the request is a cache hit versus a cache miss, we make the reasonable assumption that the user can reliably infer if its request to download an object from a server resulted in a cache hit or a cache miss immediately after the download completes.

The objective of client-side server selection is for each user  $u \in U$  to independently select a server  $s \in S$  using only the performance feedback obtained on whether each request was a hit or a miss. Let the hit rate function  $H(u, s, t)$  denote the probability of user  $u$  receiving a hit from server  $s \in S_u$  at time  $t$ . We define the system-wide performance measure  $H(t)$ , as the best hit rate obtained by the worst user at time  $t$ ,

$$H(t) \triangleq \min_{u \in U} \max_{s \in S_u} H(u, s, t), \quad (2)$$

a.k.a. the *minmax hit rate*. Our goal is to maximize  $H(t)$ . In the rest of the section, we describe a simple canonical “Go-With-The-Winner” algorithm for server selection and show that it converges quickly to an optimal state, with high probability.

*Note:* Our formulation is intentionally simple so that it could model a variety of other situations in web content delivery. For instance, a single server could in fact model a cluster of front-end servers that share a single backend object cache. A single object could in fact model a bucket of objects that cached together as is often done in a CDN context [17].

### 2.2 The *GoWithTheWinner* Algorithm

After each user  $u \in U$  has picked a content item  $a_u \in A$  using the power law distribution described in Equation 1, algorithm *GoWithTheWinner* described below is executed independently by each user  $u \in U$  to select a server that’s likely to have the content. In this algorithm, each user locally executes a simple “Go-With-The-Winner” strategy of trying out  $\sigma$  randomly chosen candidate servers initially. Then, using the past hit rate over a time window of length  $\tau$

```

1 Each user  $u$  independently chooses a random subset
   $S_u \subseteq S$  of candidate servers such that  $|S_u| = \sigma$  and
  does the following.
2 for each  $s \in S_u$  do
3   | set  $\mathbf{h}^s \leftarrow (h_1^s, h_2^s, \dots, h_\tau^s) = \mathbf{0}$ ;
4 end
5 for each arrival of request do
6   | set  $t$  to the current time;
7   | Request content  $a_u$  from all servers  $s \in S_u$ ;
8   | for each server  $s \in S_u$  do
9     |  $h_i^s \leftarrow h_{i-1}^s, 2 \leq i \leq \tau$ ;
10    |  $h_1^s \leftarrow$  if hit;  $h_1^s \leftarrow 0$ , if miss;
11    | compute hit rate  $H_\tau(u, s, t) \leftarrow (\sum_{i=1}^\tau h_i^s)/\tau$ ;
12    | if  $H_\tau(u, s, t) = 100\%$  then
13      | decide on server  $s$  by setting  $S_u \leftarrow \{s\}$ ;
14      | return;
15    | end
16  | end
17 end

```

**Algorithm 1:** GoWithTheWinner

as feedback, each user independently either chooses to continue with all the servers in  $S_u$  or decides on a single server that provided the best performance. If multiple servers provided a 100% hit rate in line 8 of the algorithm, the user decides to use the first one found.

### 2.3 Analysis of Algorithm MaxHitRate

Here we rigorously analyze the case where  $n_u = n_c = n_s = n$  and experimentally explore other variants where  $n_c$  and  $n_u$  are larger than  $n_s$  in Section 2.4 and 4. Let  $H(t)$  be as defined in (2). If  $\sigma \geq 2$ , we show that with high probability  $H(t) = 100\%$ , for all  $t \geq T$ , where  $T = O(\frac{\kappa}{\log(\kappa+1)}(\log n)^{\kappa+1} \log \log n)$ . That is, the algorithm converges quickly with high probability to an optimal state where *every* user has decided on a single server that provides a 100% hit rate, and *every* server has the content requested by its users.

*Definitions.* A server  $s$  is said to be *overbooked* at some time  $t$  if users request more than  $\kappa$  distinct content items from server  $s$ , where  $\kappa$  is the number of content items a server can hold. Note that a server may have more than  $\kappa$  users and not be overbooked, provided the users collectively request a set of  $\kappa$  or fewer content items. Also, note that a server that is overbooked at time  $t$  is overbooked at every  $t' \leq t$  since the number of users requesting a server can only remain the same or decrease with time. Finally, a user  $u$  is said to be *undecided* at time  $t$  if  $|S_u| > 1$  and is said to be *decided* if it has settled on a single server to serve its content and  $|S_u| = 1$ . Note that each user starts out undecided at time zero, then decides on a server at some time  $t$  and remains decided in all future time later than  $t$ . Users calculate the hit rates of each of the available servers based on a history record of the last  $\tau$  requests, where  $\tau$  is called the sliding window size.

**LEMMA 1.** *If the sliding window size  $\tau = \Theta(\log^{\kappa+1} n)$ , the probability that some user  $u \in U$  decides on an overbooked server  $s \in S_u$  upon any request arrival is at most  $1/n^{\Omega(1)}$ .*

**PROOF.** If user  $u$  decides on server  $s$  then the current request together with the previous  $\tau - 1$  requests are all hits. Let  $H_k, k = 1, 2, \dots, \tau$  be Bernoulli random variables,

s.t.  $H_k = 1$  if the most recent  $k$ -th request of  $u$  is a hit and  $H_k = 0$  if it is a miss. To prove Lemma 8 we need to show

$$\text{Prob}(\cap_{k=1}^\tau (H_k = 1)) \leq n^{-\Omega(1)}. \quad (3)$$

Let  $t_0$  be the time a request for content  $a$  from  $u$  is generated and appears at server  $s$ . Let  $t_0 - \Delta$  be the time that the last request for  $a$  arrives at  $s$ . Let  $H$  be an indicator variable so that  $H = 1$  if the request at  $t_0$  resulted in a hit and  $H = 0$  if resulted in a miss. Let  $A_s = \{a_1, a_2, \dots, a_M\}$  be the set of different content items requested at  $s$ , where  $M > \kappa$ . Let  $N_i$  be the number of users requesting  $a_i$  from  $s$ . WLOG, let  $a_1 = a$  be the content that  $u$  requests.  $\Delta$  is an exponential random variable,  $\Delta \sim \text{Exp}(N\lambda)$ , where  $N = N_1$  is the number of users requesting  $a$  at server  $s$ . Let  $X_i, i = 2, 3, \dots, M$  be an indicator that a request for  $a_i$  arrives at  $s$  during time interval  $(t_0 - \Delta, t_0)$ ,  $X_i \sim \text{Bernoulli}(1 - e^{-N_i\lambda\Delta})$ . Thus, random variable  $Y = \sum_{i=2}^M X_i$  is the number of requests for different content items that arrive in the time interval. With the server running LRU replacement policy,

$$\text{Prob}(H = 0) = \text{Prob}(Y \geq \kappa), \quad (4)$$

because more than  $\kappa$  different requests other than  $a$  must have arrived for content  $a$  to be swapped out of the server. (4) shows that  $H$  only depends on the arrival of other requests, which means events  $H_k, k = 1, 2, \dots, \tau$  are mutually independent. Furthermore<sup>1</sup>,

$$Y = \sum_{i=2}^M X_i \geq_d \sum_{i=2}^M X',$$

where  $X' \sim \text{Bernoulli}(1 - e^{-\lambda\Delta})$ . Furthermore because  $M \geq \kappa + 1$ ,

$$Y \geq_d Z,$$

where  $Z \sim \text{Binomial}(\kappa, (1 - e^{-\lambda\Delta}))$ .

Thus, we have

$$\begin{aligned}
\text{Prob}(Y \geq \kappa) &\geq \text{Prob}(Z \geq \kappa) \\
&= \int_0^\infty \text{Prob}(Z \geq \kappa | \Delta = t) f_\Delta(t) dt \\
&= \int_0^\infty (1 - e^{-\lambda t})^\kappa N\lambda e^{-N\lambda t} dt \\
&= \frac{N!\kappa!}{(N + \kappa)!} \\
&\geq (N + \kappa)^{-\kappa},
\end{aligned}$$

where  $f_\Delta(t)$  is the probability density function of  $\Delta$ .

Note that  $N$  is the number of users requesting  $a$  at server  $s$ , and is bounded by  $N = O(\frac{\log n}{\log \log n})$ , with high probability [19].

Now, we can finally prove (3). Let  $c'$  be an appropriate constant,

$$\begin{aligned}
\text{Prob}(\cap_{k=1}^\tau (H_k = 1)) &= \text{Prob}(H = 1)^\tau = (1 - \text{Prob}(H = 0))^\tau \\
&= (1 - \text{Prob}(Y \geq \kappa))^\tau \\
&\leq (1 - (N + \kappa)^{-\kappa})^\tau \\
&\leq (1 - (c' \frac{\log n}{\log \log n} + \kappa)^{-\kappa})^\tau,
\end{aligned}$$

<sup>1</sup>random variables  $U \geq_d V$  if  $\text{Prob}(U > x) \geq \text{Prob}(V > x)$  for all  $x$ .

which is  $n^{-\Omega(1)}$  when  $\tau = \Theta(\log^{\kappa+1} n)$ .  $\square$

By bounding the time for  $\tau$  requests to arrive at user  $u$ , we have the following,

LEMMA 2. *If user  $u$  is not decided with server  $s \in S_u$  at time  $t$ , then the server is overbooked at time  $t - \delta$  for  $\delta = \frac{\tau+1}{\lambda} c_0$  where  $c_0 > 1$  is a constant, with high probability.*

PROOF. Let random variable  $N_\delta$  be the number of requests from  $u$  during time  $(t - \delta, t)$ ,  $N_\delta \sim \text{Poisson}(\lambda\delta)$ . A bound on the tail probability of Poisson random variables is developed in [15] as

$$\text{Prob}(X \leq x) \leq \frac{e^{-\lambda'} (e\lambda')^x}{x^x},$$

where  $X \sim \text{Poisson}(\lambda')$  and  $x < \lambda'$ .

Based on that we can show there are at least  $\tau+1$  requests during  $(t - \delta, t)$  w.h.p. as the following,

$$\begin{aligned} \text{Prob}(N_\delta < \tau + 1) &\leq e^{-\lambda\delta} \frac{(e\lambda\delta)^{\tau+1}}{(\tau+1)^{\tau+1}} = e^{-(\tau+1)c_0} (ec_0)^{(\tau+1)} \\ &= e^{-(\tau+1)(c_0-1)} c_0^{(\tau+1)} \\ &= n^{-\frac{(\tau+1)}{\log n} (c_0-1-\log c_0)} \\ &= n^{-\Theta(\log^\kappa n)}, \end{aligned}$$

as  $c_0 > 1$  and  $\tau = \Theta(\log^{\kappa+1} n)$ . Thus, w.h.p. no less than  $\tau+1$  requests arrives at  $u$ . And because user is not decided at time  $t$  we know that with high probability, at least one of previous  $\tau$  requests receives a miss, which mean between the previous  $(\tau+1)$ -th request and the miss, there are  $\kappa$  different other requests arrive at the server. Thus server  $s$  is *overbooked* at the time the previous  $(\tau+1)$ -th request arrives, which with high probability is no earlier than  $t - \delta$ .  $\square$

Based on Lemmas 8 and 2, we can then establish the following theorem about the performance of Algorithm *GoWithTheWinner*.

THEOREM 3. *With probability at least  $1 - \frac{1}{n^{\Omega(1)}}$ , the min-max hit rate  $H(t) = 100\%$  for all  $t \geq T$ , provided  $\sigma \geq 2$  and  $T = O(\frac{\kappa}{\log(\kappa+1)} (\log n)^{\kappa+1} \log \log n)$ . That is, with high probability, algorithm *GoWithTheWinner* converges by time  $T$  to an optimal state where each user  $u \in U$  has decided on a server  $s \in S$  that serves it content with a 100% hit rate.*

This is the main result for the performance analysis of the algorithm. Due to space limit, please see appendix for detailed proof of this theorem.

Are two or more random choices necessary for all users to receive a 100% hit rate? Analogous to the “power of two choices” in the balls-into-bins context [14], we show that two or more choices are required for good performance.

THEOREM 4. *For any fixed constants  $0 \leq \alpha < 1$  and  $\kappa \geq 1$ , when algorithm *MaxHitRate* uses one random choice for each user ( $\sigma = 1$ ), the minmax hit rate  $H(t) = o(1)$ , with high probability, i.e.,  $H(t)$  tends to zero as  $n$  tends to infinity, with high probability.*

Please see appendix for the proof.

## 2.4 The case when $n_u = n_s \times \log n_s$

Now we analyze the case that there are much more users than the number of servers, which is often the case in reality. Assume  $n_u = n_s \log n_s$  and  $\kappa = \log n_s$ , we have the following results,

LEMMA 5. *When  $n_u = n_s \log n_s$ , with probability at least  $1 - \frac{1}{n_s^{\Omega(1)}}$ , the maximum load (number of incoming servers) over all servers is  $O(\sigma \frac{n_u}{n_s})$ . Furthermore, if  $\kappa = \log n_s$ , all users have 100% hit rate.*

Lemma 5 implies that when  $n_u = n_s \log n_s$  all the servers have balanced load of  $\sigma \frac{n_u}{n_s}$ , thus we don’t need more server selection mechanism for load balancing other than just letting all users randomly choose the server. And in this case, it’s not beneficial to let users start with more than 1 randomly selected servers, because with  $\sigma = 1$  the load on all servers are balanced already. Thus, as long as we have feasible server capacity  $\kappa = \frac{n_u}{n_s}$ , all the users will have enough resources from the server and have 100% hit rate by randomly select 1 server.

The number of content items  $n_c$  here doesn’t not affect the result of load balancing. Actually, the result stays the same when  $n_c \geq n_u$ . And when the number of content items is much smaller than number of users,  $n_c \ll n_u$ , the cache size can become smaller ( $\kappa < \frac{n_u}{n_s}$ ) because the number of distinct requests at each server becomes smaller.

## 3. BITRATE MAXIMIZATION FOR VIDEO CONTENT

In video streaming, a key performance metric is the *bitrate* at which an user can download the video stream. If the server is unable to provide the required bitrate to the user, the video will freeze frequently resulting in an inferior viewing experience and reduced user engagement [12]. For simplicity, we model the server’s bandwidth capacity that is often the critical bottleneck resource, while leaving other factors that could influence video performance such as the server-to-user connection and the server’s cache<sup>2</sup> for future work.

### 3.1 Problem formulation

The bitrate required to play a stream without freezes is often the encoded bitrate of the stream. For simplicity, we assume that each user requires a bitrate of 1 unit for playing its video and each server has the capacity to serve  $\kappa$  units in aggregate. And we assume each server evenly divides its available bitrate capacity among all users who keeps a streaming connection with it. We make the reasonable assumption that each user can compute the bitrate that it receives from its chosen candidate servers and that this bitrate is used as the performance feedback (cf. Figure 1).

Different from the delivering web content, where users make repetitive requests to the same website with Poisson processes, we consider users for video streaming have persistent connection with the server. We use a discrete time model in this case as compared to web content delivery

<sup>2</sup>Unlike the web, cache hit rate is a less critical determinant of video performance. Videos are cached in chunks by the server. The next chunk is often prefetched from origin if it is not in cache, even while the current chunk is being played by the user, so as to hide the origin-to-server latency.

where everything is in continuous time. We assume after each time unit, the users look at the bit rate provided by each of the available servers and then make decisions according to the performance (measured by bit rate). The goal of each user is to find a server who can provide the required bitrate of 1 unit for viewing the video.

### 3.2 Algorithm MaxBitRate

After each user  $u \in U$  has picked a video object  $c_u \in C$  using the power law distribution described in Equation 1, Algorithm MaxBitRate described below is executed independently by each user  $u \in U$ , in discrete time steps.

1. Choose a random subset of candidate servers  $S_u \subseteq S$  such that  $|S_u| = \sigma$ .
2. At each time step  $t \geq 0$ , do the following:
  - (a) Request the video content from all servers  $s \in S_u$ .
  - (b) For each server  $s \in S_u$ , compute  $B(u, s, t) \triangleq$  bitrate provided by server  $s$  to user  $u$  in the current time step.
  - (c) If there exists a server  $s \in S_u$  such that  $B(u, s, t) = 1$ , then *decide* on server  $s$  by setting  $S_u \leftarrow \{s\}$ .

Note that the users are executing a simple strategy of trying out  $\sigma$  randomly chosen servers initially. Then, using the bitrate received in the current time step as feedback, each user independently narrows its choice of servers to a single server that provided the required unit bitrate. If multiple servers provided the required bitrate, the user decides to use an arbitrary one. Further, note that a user  $u$  downloading from a server  $s$  at time  $t$  knows immediately whether or not the server is overloaded, since server  $s$  is overloaded iff user  $u$  received a bitrate of less than 1 unit from the server, i.e.,  $B(u, s, t) < 1$ . This is a point of simplification in relation to the more complex situation for hit rate maximization where any single cache hit is not indicative of a non-overloaded server and a historical average of hit rates over a large enough time window  $\tau$  is required as a probabilistic indicator of server overload. And furthermore, this simplification yields both faster convergence to an optimal state in  $T = O(\log \log n / \log(\kappa + 1))$  steps and a much simpler proof of that convergence.

### 3.3 Analysis of Algorithm MaxBitRate

As before, we rigorously analyze the case where  $n_u = n_s = n_c = n$ . Let the *minmax* bitrate  $B(t)$  be the best bitrate obtained by the worst user at time  $t$ , i.e.,

$$B(t) \triangleq \min_{u \in U} \max_{s \in S} B(u, s, t).$$

**THEOREM 6.** *When  $\sigma \geq 2$ , the minmax bitrate converges to  $B(t) = 1$  unit, for all  $t \geq T$ , within time  $T = O(\log \log n / \log(\kappa + 1))$ , with high probability. When  $\sigma = 1$  on the other hand, the minmax bitrate  $B(t) = O(\kappa \log \log n / \log n)$ , with high probability. In particular, when  $\sigma = 1$  and the cache size  $\kappa$  is  $o(\log n / \log \log n)$ , including the case when  $\kappa$  is a fixed constant,  $B(t)$  tends to zero as  $n$  tends to infinity, with high probability.*

Please refer to appendix for the proof.

## 4. EMPIRICAL EVALUATION

We empirically study our algorithm *GoWithTheWinner* by building a simulator. Each user is implemented as a Poisson arrival sequence with unit rate. We use  $n_u = 1000$  users. To simulate varying numbers of servers, users, and applications, we also varied  $n_s$  and  $n_a$  such that  $1 \leq n_u/n_c, n_u/n_s \leq 100$ . We also simulate a range of values for the spread  $1 \leq \sigma \leq 6$ , and sliding window size  $1 \leq \tau \leq 20$ . Each server implements an LRU application replacement policy of size  $\kappa \geq 2$ . The applications are requested by users using the power law distribution of Equation 1 with  $\alpha = 0.65$  to model realistic content popularity [4] [10]. However, we also vary  $\alpha$  from 0 (uniform distribution) to 1.5 in some of our simulations.

The system is said to have *converged* when all users have decided on a single server from their set of candidate servers. There are two complementary metrics that relate to convergence. *Failure rate* is the probability that the system converged to a non-optimal state where there exists servers that are overbooked, resulting in some users incurring application misses after convergence occurred. The failure rate is measured by performing the simulation multiple times and assessing the goodness of the converged state. *Convergence time* is the time it takes for the system to converge provided that it converged to an optimal state.

### 4.1 Speed of convergence

Figure 2 shows how the fraction of undecided users decreases over time till it reaches zero, resulting in convergence. Note that users do not decide in the first  $\tau$  steps, since they must wait at least that long to accumulate a window of  $\tau$  application hits. However, once the first  $\tau$  steps complete, the decrease in the number of undecided users is fast as users discover that at least one of their two randomly chosen candidate servers have less load. The rate of decrease in undecided users slows down again towards the end, as pockets of users who experience cache contention in *both* of their server choices require multiple iterations to resolve.

In this simulation, we keep the number of users  $n_u = 1000$  but vary the number of servers  $n_s$  to achieve different values for  $n_u/n_s$ . Note that for a fair comparison, we keep the system-wide load the same. Load  $l$  is a measure of cache contention in the network and is naturally defined as the ratio of the numbers of users in the system and total serving capacity that is available in the system. That is,  $l \triangleq n_u / (\kappa \cdot n_s)$ . For all three setting of Figure 2, we keep load  $l = 0.5$ . The figure shows that with fewer (but larger) servers ( $n_u/n_s$  is larger) the convergence time is faster, because having server capacity in a few larger servers provides a larger application hit rate than having the same capacity in several smaller servers. Similar performance gains are also found in the context of web caching and parallel jobs scheduling [18]. The convergence times are plotted explicitly in Figure 4 for a greater range of user-to-server ratios. As  $n_u/n_s$  increases from 1 to 40, convergence time decreases. The decrease in convergence times are not significant beyond  $n_u/n_s \geq 40$ .

### 4.2 Impact of sliding window $\tau$

The sliding window  $\tau$  is the number of recent requests used by algorithm *GoWithTheWinner* to estimate the hit rate. As shown in Figure 3, there is a natural tradeoff between convergence time and failure rate. When  $\tau$  increases,




Figure 2: The figures show the percentage of undecided users for a typical power law distribution ( $\alpha = 0.65$ ) with spread  $\sigma = 2$  and  $n_u = 1000$ . Note that the undecided users decrease with time in all cases, but the convergence is faster when we use fewer but larger servers by setting  $n_u/n_s$  to be larger. Also, the smaller values of the look-ahead window  $\tau$  result in faster convergence.




Figure 3: Generally, as  $\tau$  increases, convergence time increases but failure rate decreases. It is also true for larger servers ( $n_u/n_s = 20$ ), only the failure has gone to zero for all investigated sliding window size  $\tau$ .

the users take longer to converge, as they require a 100% hit rate in a larger sliding window. However, waiting for a longer period also makes their decisions more robust. That is, a user is less likely to choose an overbooked server, since an overbooked server is less likely to provide a string of  $\tau$  application hits for large  $\tau$ . In our simulations with many smaller caches ( $n_u/n_s = 1$ ), when  $\tau \leq 4$ , users made quick choices based on a smaller sliding window. But, this resulted in the system converging to a non-optimal state 100% of the time. As  $\tau$  further increases, the failure rate decreased. The value of  $\tau = 11$  is a suitable sweet spot as it results in the smallest convergence time for a zero failure rate. However, for fewer but larger servers ( $n_s/n_u = 20$ ), all selections of window size  $\tau$  (thus the small values like  $\tau = 5$ ) yielded a 0% failure rate, while the convergence time still increases as the window size gets larger.

### 4.3 Impact of spread $\sigma$

As shown in Theorems 3 and 4, a spread of  $\sigma \geq 2$  is required for the system to converge to an optimal solution, while a spread of  $\sigma = 1$  is insufficient. As predicted by our analysis, our simulations did not converge to an optimal state with  $\sigma = 1$ . Figure 5 shows the convergence time as a function of spread, for  $\sigma \geq 2$ .

As  $\sigma$  increases, there are two opposing factors that impact the convergence time. The first factor is that as  $\sigma$  increases, each user has more choices and the user is more likely to find a suitable server with less load. On the other hand,

an increase in  $\sigma$  also increases the total number of initial requests in the system that equals  $\sigma n_u$ . Thus, the system starts out in a state where servers have greater average load when  $\sigma$  is larger. These opposing forces result in a very small incremental benefit when using  $\sigma = 3$  instead of 2, though the higher values of  $\sigma > 3$  showed no benefit as convergence time increases with  $\sigma$  increases.

We established the “power of two random choices” phenomenon where two or more random server choices yield superior results to having just one. It is intriguing to ask *what percentage of users need two choices* to reap the benefits of multiple choices? Consider a mix of users, some with two random choices and others with just one. Let  $\sigma_{avg}$ ,  $1 \leq \sigma_{avg} \leq 2$ , denote the average value of the spread among the users.

In Figure 6, we show different order statistics of the hit rate as a function of  $\sigma_{avg}$ . Specifically, we plot the minimum value, 1<sup>st</sup>-percentile, 5<sup>th</sup>-percentile and the median (50<sup>th</sup>-percentile) of the hit rates of the users after running the system for a long enough period of 200 time units. As our theory predicts, when  $\sigma_{avg} = 2$ , the minimum and all the order statistics converge to 100%, as all users converge to a 100% hit rate. Further, if we are interested in only the median user, any value of the spread is sufficient to guarantee that 50% of the users obtain a 100% hit rate. Perhaps the most interesting phenomena is that if  $\sigma_{avg} = 1.7$ , i.e., 70% of the users have two choices and the rest have one choice, the 5<sup>th</sup>-percentile converges to 100%, i.e., all but 5%






Figure 4: As  $n_u/n_s$  increases fewer servers with larger capacity are used and convergence time decreases. The decrease is less pronounced beyond  $n_u/n_s \geq 40$  under this setting ( $\alpha = 0.65$ ,  $\sigma = 2$ ,  $\tau = 20$ ).




Figure 6: Order statistics of the hit rate of the user population. ( $\alpha = 0.65$ ,  $n_u/n_s = 1$ ,  $\tau = 10$ ,  $\kappa = 2$ .)

of the users experience a 100% hit rate. For a higher value of  $\sigma_{avg} = 1.9$ , the 1<sup>st</sup>-percentile converges to 100%, i.e., all but the 1% of the users experience a 100% hit rate. This result shows that our algorithm still provides benefits even if only *some* users have multiple random choices of servers available to them.

#### 4.4 Impact of demand distribution

We now study how hit rate changes with the exponent  $\alpha$  in the power law distribution of Equation 1. Note that the distribution is uniform when  $\alpha = 0$  and is the harmonic distribution when  $\alpha = 1$ . As  $\alpha$  increases, since the tails fall as a power of  $\alpha$ , the distribution gets more and more skewed towards applications with a smaller rank. In Figure 7, we plot the minmax hitrate over time for different  $\alpha$ , where we see that a larger  $\alpha$  leads to faster convergence. The reason is that as the popularity distribution gets more skewed, a larger fraction of users can share the same VMs for popular applications, leading to better hit rate and faster convergence. Thus, the uniform application popularity distribution ( $\alpha = 0$ ) is the worst case and the algorithm converges faster for the distributions that tend to occur more commonly in practice. Providing theoretical support for this empirical result by analyzing the convergence time to show faster convergence for larger  $\alpha$  is a topic for future work.

### 5. RELATED WORK

Server selection algorithms have a rich history of both research and actual implementations over the past two decades. Several server selection algorithms have been proposed and empirically evaluated, including client-side algorithms that use historical performance feedback using probes [8, 6]. Server selection has also been studied in a variety of contexts, such




Figure 5: There is a very small incremental benefit in using  $\sigma = 3$  instead of 2, though higher values of  $\sigma > 3$  only increased the convergence time. ( $\alpha = 0.65$ ,  $n_u/n_s = 1$ ,  $\tau = 20$ ,  $\kappa = 2$ .)




Figure 7: Minmax hitrate versus time for different power law distributions.

as the web [6, 20], video streaming[21], and cloud services[22]. Our work is distinguished from the prior literature in that we theoretically model the “Go-With-The-Winner” paradigm that is common to many proposed and implemented client-side server selection algorithms. Our work is the first formal study of the efficacy and convergence of such algorithms.

In terms of analytical techniques, our work is closely related to prior work on balls-into-bins games where the witness tree technique was first utilized [14]. Witness trees were subsequently used to analyze load balancing algorithms, and circuit-switching algorithms [5]. However, our setting involves additional complexity requiring novel analysis due to the fact that users can share a single cached copy of an object and the hitrate feedback is only a probabilistic indicator of server overload. Also, our work shows that the “power of two random choices” phenomenon applies in the context of content delivery, a phenomenon known to hold in other contexts such as balls-into-bins, load balancing, relay allocation for services like Skype [16], and circuit switching in interconnection networks [14].

### 6. CONCLUSION

Our work constitutes the first formal study of the simple “Go-With-The-Winner” paradigm in the context of web and video content delivery. For web (resp., video) delivery, we proposed a simple algorithm where each user randomly chooses two or more candidate servers and selects the server that provided the best hitrate (resp., bitrate). We proved that the algorithm converges quickly to an optimal state where all users receive the best hitrate (resp., bitrate) and no server is overloaded, with high probability. While we make some assumptions to simplify the theoretical analysis, our simulations evaluate a broader setting that



incorporates a range of values for  $\tau$  and  $\sigma$ , varying content popularity distributions, differing load conditions, and situations where only some users have multiple server choices. Taken together, our work establishes that the simple “Go-With-The-Winner” paradigm can provide algorithms that converge quickly to an optimal solution, given a sufficient number of random choices and a sufficiently (but not perfectly) accurate performance feedback.

## 7. REFERENCES

- [1] V. K. Adhikari, Y. Guo, F. Hao, M. Varvello, V. Hilt, M. Steiner, and Z.-L. Zhang. Unreeling netflix: Understanding and improving multi-cdn movie delivery. In *INFOCOM, 2012 Proceedings IEEE*, pages 1620–1628. IEEE, 2012.
- [2] Akamai. Akamai download manager. 2013. [http://www.akamai.com/html/solutions/downloadmanager\\_overview.html](http://www.akamai.com/html/solutions/downloadmanager_overview.html).
- [3] M. Bellare and T. Kohno. Hash function balance and its impact on birthday attacks. In *Advances in Cryptology-Eurocrypt 2004*, pages 401–418. Springer, 2004.
- [4] L. Breslau, P. Cue, P. Cao, L. Fan, G. Phillips, and S. Shenker. Web caching and zipf-like distributions: Evidence and implications. In *In INFOCOM*, pages 126–134, 1999.
- [5] R. Cole, B. M. Maggs, M. Mitzenmacher, A. W. Richa, K. Schröder, R. K. Sitaraman, B. Vöcking, et al. Randomized protocols for low-congestion circuit routing in multistage interconnection networks. In *Proceedings of the thirtieth annual ACM symposium on Theory of computing*, pages 378–388. ACM, 1998.
- [6] M. E. Crovella and R. L. Carter. Dynamic server selection in the internet. Technical report, Boston University Computer Science Department, 1995.
- [7] J. Dilley, B. Maggs, J. Parikh, H. Prokop, R. Sitaraman, and B. Weihl. Globally distributed content delivery. *Internet Computing, IEEE*, 6(5):50–58, 2002.
- [8] S. G. Dykes, K. A. Robbins, and C. L. Jeffery. An empirical evaluation of client-side server selection algorithms. In *INFOCOM 2000. Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, volume 3, pages 1361–1370. IEEE, 2000.
- [9] C. Fricker, P. Robert, and J. Roberts. A versatile and accurate approximation for lru cache performance. In *Proceedings of the 24th International Teletraffic Congress, ITC '12*, pages 8:1–8:8. International Teletraffic Congress, 2012.
- [10] P. Gill, M. Arlitt, Z. Li, and A. Mahanti. Youtube traffic characterization: A view from the edge. In *Proceedings of the 7th ACM SIGCOMM Conference on Internet Measurement, IMC '07*, pages 15–28, New York, NY, USA, 2007. ACM.
- [11] P. Key, L. Massoulie, and P. Towsley. Path selection and multipath congestion control. *INFOCOM 2007. 26th IEEE International Conference on Computer Communications. IEEE*, pages 143–151, may 2007.
- [12] S. S. Krishnan and R. K. Sitaraman. Video stream quality impacts viewer behavior: inferring causality using quasi-experimental designs. In *Proceedings of the 2012 ACM conference on Internet measurement conference*, pages 211–224. ACM, 2012.
- [13] C. Liu and P. Albitz. *DNS and Bind*. O’Reilly Media, Inc., 2009.
- [14] M. Mitzenmacher, A. W. Richa, and R. Sitaraman. The power of two random choices: A survey of techniques and results. *COMBINATORIAL OPTIMIZATION-DORDRECHT*, 9(1):255–304, 2001.

- [15] M. Mitzenmacher and E. Upfal. *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*. Cambridge University Press, New York, NY, USA, 2005.
- [16] H. X. Nguyen, D. R. Figueiredo, M. Grossglauser, and P. Thiran. Balanced relay allocation on heterogeneous unstructured overlays. In *INFOCOM*, pages 126–130, 2008.
- [17] E. Nygren, R. K. Sitaraman, and J. Sun. The akamai network: a platform for high-performance internet applications. *ACM SIGOPS Operating Systems Review*, 44(3):2–19, 2010.
- [18] K. Ousterhout, P. Wendell, M. Zaharia, and I. Stoica. Sparrow: Distributed, low latency scheduling. In *Proceedings of the Twenty-Fourth ACM Symposium on Operating Systems Principles*. ACM.
- [19] M. Raab and A. Steger. Balls into bins—A simple and tight analysis. In *Randomization and Approximation Techniques in Computer Science*, pages 159–170. Springer, 1998.
- [20] M. Sayal, Y. Breitbart, P. Scheuermann, and R. Vingralek. Selection algorithms for replicated web servers. *ACM SIGMETRICS Performance Evaluation Review*, 26(3):44–50, 1998.
- [21] R. Torres, A. Finamore, J. R. Kim, M. Mellia, M. M. Munafo, and S. Rao. Dissecting video server selection strategies in the youtube cdn. In *Distributed Computing Systems (ICDCS), 2011 31st International Conference on*, pages 248–257. IEEE, 2011.
- [22] P. Wendell, J. W. Jiang, M. J. Freedman, and J. Rexford. Donar: decentralized server selection for cloud services. In *ACM SIGCOMM Computer Communication Review*, volume 40, pages 231–242. ACM, 2010.

## APPENDIX

### 1 Detailed Proof of Theorem 3

PROOF. For simplicity, we prove the situation where  $\sigma = 2$ , i.e., each user initially chooses two random candidate servers in step 1 of the algorithm. The case where  $\sigma > 2$  is analogous. Wlog, we also assume  $\kappa$  is at most  $O(\log n / \log \log n)$ , which includes the interesting case of  $\kappa$  equal to a constant. When the server capacity is larger, i.e., if  $\kappa = \Omega(\log n / \log \log n)$ , there will be no overbooked servers with high probability and the theorem holds trivially. This observation follows from a well-known result that if  $n$  balls (i.e., users) uniformly and randomly select  $k$  out of  $n$  bins (i.e. servers), then the maximum number of users that select a server is  $O(\log n / \log \log n)$  with high probability, when  $k$  is a fixed constant [19].

In contradiction to the theorem, suppose some user  $u$  has not decided on a server by time  $T$ . We construct a “witness tree<sup>3</sup>” of degree  $\kappa + 1$  and depth at least  $\rho$ , where  $\rho = T/\delta = \kappa \log \log n / \log(\kappa + 1)$ . Each node of the witness tree is a server. Each edge of the witness tree is a user whose two nodes correspond to the two servers chosen by that user. We show that the existence of an undecided user in time step  $T$  is unlikely by enumerating all possible witness trees

and showing that the occurrence of any such witness tree is unlikely. The proof proceeds in the following three steps.

**(1) Constructing a witness tree.** If algorithm MaxHitRate has not converged to the optimal state at time  $T$ , then there exists a user (say  $u_1$ ) and a server  $s$  such that  $H_\tau(u_1, s, T) < 100\%$ , since user  $u_1$  has not yet found a server with a 100% hit rate. We make server  $s$  the root of the witness tree.

We find children for the root  $s$  to extend the witness tree as follows. Since  $H_\tau(u_1, s, T) < 100\%$ , by Lemma 2 we know server  $s$  is overbooked at time  $t' = t - \delta$ , i.e., there are at least  $\kappa + 1$  users requesting server  $s$  for  $\kappa + 1$  distinct applications at time  $t'$ . Let  $u_1, \dots, u_{\kappa+1}$  be the users who sent requests to server  $s$  at time  $t'$ . Wlog, assume that the users  $\{u_i\}$  are ordered in ascending order of their IDs. By Lemma 8, we know that the probability of a user deciding on an overbooked server is small, i.e., at most  $1/n^{\Omega(1)}$ . Thus, with high probability, users  $u_1, \dots, u_{\kappa+1}$  are undecided at time  $t'$  since server  $s$  is overbooked. Let  $s_i$  be the other server choice associated with user  $u_i$  (one of the choices is server  $s$ ). We extend the witness tree by creating  $\kappa + 1$  children for the root  $s$ , one corresponding to each server  $s_i$ . Note that for each of the servers  $s_i$  we know that  $H(u_i, s_i, t') < 100\%$ , since otherwise user  $u_i$  would have decided on server  $s_i$  in time step  $t'$ . Thus, analogous to how we found children for  $s$ , we can recursively find  $\kappa + 1$  children for each of the servers  $s_i$  and grow the witness tree to an additional level.

Observe that to add an additional level of the witness tree we went from server  $s$  at time  $T$  to servers  $s_i$  at time  $t'$ , i.e., we went back in time by an amount of  $T - t' \leq \delta$ . If we continue the same process, we can construct a witness tree that is a  $(\kappa + 1)$ -ary tree of depth  $T/\delta = \rho$ .

**(2) Pruning the witness tree.** If the nodes of the witness tree are guaranteed to represent distinct servers, proving our probabilistic bound is relatively easy. The reason is that if the servers are unique then the users that represent edges of the tree are unique as well. Therefore the probabilistic choices that each user makes is independent, making it easy to evaluate the probability of occurrence of the tree. However, it may not be the case that the servers in the witness tree constructed above are unique, leading to dependent choices that are hard to resolve. Thus, we create a *pruned witness tree* by removing repeated servers from the original (unpruned) witness tree.

We prune the witness tree by visiting the nodes of the witness tree iteratively in breadth-first search order starting at the root. As we perform breadth-first search (BFS), we remove (i.e., prune) some nodes of the tree and the subtrees rooted at these nodes. What is left after this process is the pruned witness tree. We start by visiting the root. In each iteration, we visit the next node  $v$  in BFS order that has not been pruned. Let  $\beta(v)$  denote the nodes visited *before*  $v$ . If  $v$  represents a server that is different from the servers represented by nodes in  $\beta(v)$ , we do nothing. Otherwise, prune all nodes in the subtree rooted at  $v$ . Then, mark the edge from  $v$  to its parent as a *pruning edge*. (Note that the pruning edges are not part of the pruned witness tree.) The procedure continues until either no more nodes remain to be visited or there are  $\kappa + 1$  pruning edges. In the latter case, we apply a final pruning by removing all nodes that are yet to be visited, though this step does not produce any more pruning edges. This process results in a pruned witness and a set of  $p$  (say) pruning edges.

Note that each pruning edge corresponds to a user who

<sup>3</sup>A witness tree is so called as it bears witness to the occurrence of an event such as a user being undecided.

we will call a *pruned user*. We now make a pass through the pruning edges to select a set  $P$  of unique pruned users. Initially,  $P$  is set to  $\emptyset$ . We visit the pruning edges in BFS order and for each pruning edge  $(u, v)$  we add the user corresponding to  $(u, v)$  to  $P$ , if this user is distinct from all users currently in  $P$  and if  $|P| < \lceil p/2 \rceil$ , where  $p$  is the total number of pruning edges. We stop adding pruned users to set  $P$  when we have exactly  $\lceil p/2 \rceil$  users. Note that since a user who made server choices of  $u$  and  $v$  can appear at most twice as a pruned edge, once with  $u$  in the pruned witness tree and once with  $v$  in the pruned witness tree. Thus, we are guaranteed to find  $\lceil p/2 \rceil$  distinct pruned users.

After the pruning process, we are left with a pruned witness tree with nodes representing distinct servers and edges representing distinct users. In addition, we have a set  $P$  of  $\lceil p/2 \rceil$  distinct pruned users, where  $p$  is the number of pruning edges.

**(3) Bounding the probability of pruned witness trees.** We enumerate possible witness trees and bound their probability using the union bound. Observe that since the (unpruned) witness tree is a  $(\kappa + 1)$ -ary tree of depth  $\rho$ , the number of nodes in the witness tree is

$$m = \sum_{0 \leq i \leq \rho} (\kappa + 1)^i = \frac{(\kappa + 1)^{\rho+1} - 1}{\kappa} \leq 2 \log^2 n, \quad (5)$$

since  $\rho = 2 \log \log n / \log(\kappa + 1)$  and hence  $(\kappa + 1)^\rho = \log^2 n$ .

*Ways of choosing the shape of the pruned witness tree:* The shape of the pruned witness tree is determined by choosing the  $p$  pruning edges of the tree. The number of ways of selecting the  $p$  pruning edges is at most  $\binom{m}{p} \leq m^p$ , since there are at most  $m$  edges in the (unpruned) witness tree.

*Ways of choosing users and servers for the nodes and edges of the pruned witness tree:* The enumeration proceeds by considering the nodes in BFS order. The number of ways of choosing the server associated with the root is  $n$ . Consider the  $i^{\text{th}}$  internal node  $v_i$  of the pruned witness tree whose server has already been chosen to be  $s_i$ . Let  $v_i$  have  $\mu_i$  children. There are at most  $\binom{n}{\mu_i}$  ways of choosing distinct servers for each of the  $\mu_i$  children of  $v_i$ . Also, since there are at most  $n$  users in the system at any point in time, the number of ways to choose distinct users for the  $\mu_i$  edges incident on  $v_i$  is also at most  $\binom{n}{\mu_i}$ . There are  $\mu_i!$  ways of pairing the users and the servers. Further, the probability that a chosen user chooses server  $s_i$  corresponding to node  $v_i$  and a specific one of  $\mu_i$  servers chosen above for  $v_i$ 's children is

$$\frac{1}{\binom{n}{2}} = \frac{2}{n(n-1)},$$

since each set of two servers is equally likely to be chosen in step 1 of the algorithm. Further, note that each of the  $\mu_i$  users chose  $\mu_i$  distinct applications and let the probability of occurrence of this event be  $\text{Uniq}(n_a, \mu_i)$ . This uniqueness probability has been studied in the context of collision-resistant hashing and it is known [3] that  $\text{Uniq}(n_a, \mu_i)$  is largest when the content popularity distribution is the uniform distribution ( $\alpha = 0$ ) and progressively becomes smaller as  $\alpha$  increases. In particular,  $\text{Uniq}(n_a, \mu_i) \leq e^{-\Theta(\mu_i^2/n_a)} < 1$ . Putting it together, the number of ways of choosing a distinct server for each of the  $\mu_i$  children of  $v_i$ , choosing a distinct user for each of the  $\mu_i$  edges incident on  $v_i$ , choosing a distinct application for each user, and multiplying by the

appropriate probability is at most

$$\binom{n}{\mu_i} \cdot \binom{n}{\mu_i} \cdot \mu_i! \cdot \left( \frac{2}{n(n-1)} \right)^{\mu_i} \cdot \text{Uniq}(n_a, \mu_i) \leq \frac{2^{\mu_i}}{\mu_i!}, \quad (6)$$

provided  $\mu_i > 1$ . Let  $m'$  be the number of internal nodes  $v_i$  in the pruned witness tree such that  $\mu_i = \kappa + 1$ . Using the bound in Equation 9 for only these  $m'$  nodes, the number of ways of choosing the users and servers for the nodes and edges respectively of the pruned witness tree weighted by the probability that these choices occurred is at most

$$n \cdot (2^{\kappa+1}/(\kappa+1)!)^{m'}.$$

*Ways of choosing the pruned users in  $P$ :* Recall that there are  $\lceil p/2 \rceil$  distinct pruned users in  $P$ . The number of ways of choosing the users in  $P$  is at most  $n^{\lceil p/2 \rceil}$ , since at any time step there are at most  $n$  users in the system to choose from. Note that a pruned user has both of its server choices in the pruned witness tree. Therefore, the probability that a given user is a pruned user is at most  $m^2/n^2$ . Thus the number of choices for the  $\lceil p/2 \rceil$  pruned users in  $P$  weighted by the probability that these pruned users occurred is at most

$$n^{\lceil p/2 \rceil} \cdot (m^2/n^2)^{\lceil p/2 \rceil} \leq (m^2/n)^{\lceil p/2 \rceil}.$$

*Bringing it all together:* The probability that there exists a pruned witness tree with  $p$  pruning edges, and  $m'$  internal nodes with  $(\kappa + 1)$  children each, is at most

$$\begin{aligned} & m^p \cdot n \cdot (2^{\kappa+1}/(\kappa+1)!)^{m'} \cdot (m^2/n)^{\lceil p/2 \rceil} \\ & \leq n \cdot (2^{\kappa+1}/(\kappa+1)!)^{m'} \cdot (m^4/n)^{\lceil p/2 \rceil} \\ & \leq n \cdot (2e/(\kappa+1))^{m'(\kappa+1)} \cdot (m^4/n)^{\lceil p/2 \rceil}, \end{aligned} \quad (7)$$

since  $(\kappa + 1)! \geq ((\kappa + 1)/e)^{\kappa+1}$ . There are two possible cases depending on how the pruning process terminates. If the number of pruning edges,  $p$ , equals  $\kappa + 1$  then the third term of Equation 10 is

$$(m^4/n)^{\lceil p/2 \rceil} \leq (16 \log^8 n/n)^{\lceil (\kappa+1)/2 \rceil} \leq 1/n^{\Omega(1)},$$

using Equation 8 and assuming that cache size  $\kappa$  is at least a suitably large constant. Alternately, if the pruning process terminates with fewer than  $\kappa + 1$  pruning edges, it must be that at least one of the  $\kappa + 1$  subtrees rooted at the children of the root  $s$  of the (unpruned) witness tree has no pruning edge. Thus, the number of internal nodes  $m'$  of the pruned witness tree with  $(\kappa + 1)$  children each is bounded as follows:

$$m' = \sum_{0 \leq i < \rho-1} (\kappa + 1)^i \geq (\kappa + 1)^{\rho-2} \geq \log^2 n / (\kappa + 1)^2,$$

as  $(\kappa + 1)^\rho = \log^2 n$ . Thus, the second term of Equation 10 is

$$(2e/(\kappa + 1))^{m'(\kappa+1)} \leq (2e/(\kappa + 1))^{\log^2 n / (\kappa+1)} \leq 1/n^{\Omega(1)},$$

assuming  $\kappa > 2e - 1$  but is at most  $O(\log n / \log \log n)$ . Thus, in either case, the bound in Equation 10 is  $1/n^{\Omega(1)}$ . Further, since there are at most  $m$  values for  $p$ , the total probability of a pruned witness tree is at most  $m \cdot 1/n^{\Omega(1)}$  which is  $1/n^{\Omega(1)}$ . This completes the proof of the theorem.  $\square$

## 2 Proof of Theorem 4

**PROOF.** From the classical analysis of throwing  $n$  balls into  $n$  bins [14], we know that there exist a subset  $U' \subseteq U$

such that  $|U'| = \Theta(\log n / \log \log n)$  and all users in  $U'$  have chosen a single server  $s$ , with high probability. Now we show that some user in  $U'$  must have a small hit rate with high probability. Let  $C'$  represent the set of all objects accessed by all users in  $S'$ . The probability that  $|C'| \leq \kappa w(n)$  can be upper bounded as follows, where  $w(n)$  is an arbitrarily slowly growing function of  $n$ . The number of ways of picking  $C'$  objects from a set  $C$  of  $n$  objects is at most  $n^{|C'|}$ . The probability that a user in  $U'$  will pick an object in  $C'$  can be upper bounded by the probability that a user chooses one of the  $|C'|$  most popular objects. Thus the probability that a user in  $U'$  picks an object in  $C'$  is at most  $\mathcal{H}(|C'|, \alpha) / \mathcal{H}(n, \alpha) = \Theta((|C'|/n)^{1-\alpha})$ , where  $\mathcal{H}(i, \alpha)$  is the  $i^{\text{th}}$  generalized harmonic number and  $\mathcal{H}(i, \alpha) = \Theta(i^{1-\alpha})$ . Thus, the probability that all users in  $U'$  pick objects in  $C'$  is at most  $\Theta((|C'|/n)^{(1-\alpha)|U'|})$ . Therefore, the probability that  $|C'| \leq \kappa w(n)$  is at most

$$\begin{aligned} & n^{|C'|} \cdot \Theta((|C'|/n)^{(1-\alpha)|U'|}) \\ & \leq n^{\kappa w(n)} \cdot (\kappa w(n)/n)^{\Theta(\log n / \log \log n)} = o(1) \end{aligned}$$

Thus, probability that  $|C'| \leq \kappa w(n)$  is small and hence  $|C'| > \kappa w(n)$ , with high probability. Since the minmax hit rate  $H(t)$  is at most  $\kappa/|C'|$  which is at most  $1/w(n)$ ,  $H(t)$  tends to zero with high probability.  $\square$

### 3 Proof of Lemma 5

PROOF. We prove the lemma using Chernoff Bound. We firstly look at the load of one server. Let  $X_i$  be the indicator that user  $u_i$  selected the server we looked at, and let  $Y = \sum_{i=1}^{n_u} X_i$  be the total number of incoming users at this server. Because each user chooses  $\sigma$  servers uniformly at random, we have  $\text{Prob}(X_i = 1) = \sigma/n_s$ . Thus we have,

$$\begin{aligned} \text{Prob}\left(Y > \sigma \frac{n_u}{n_s}\right) &= \text{Prob}(Y > \sigma \log n_s) \\ &= \text{Prob}(Y \geq \sigma \log n_s + 1) \\ &= \text{Prob}\left(Y \geq \left(1 + \frac{1}{\sigma \log n_s}\right) \sigma \log n_s\right) \\ &\leq e^{-n_s \log n_s \frac{1}{3(\sigma \log n_s)^2}} \\ &= n_s^{-\frac{n_s}{3\sigma^2 \log^2(n_s)}}. \end{aligned}$$

With which we can then calculate the bound on the maximum load with union bound. Let  $Y_i$  be the number of users at server  $u_i$ , then we have

$$\begin{aligned} \text{Prob}\left(\max_i \{Y_i\} > \sigma \frac{n_u}{n_s}\right) &= \text{Prob}\left(\bigcup_{i=1}^{n_s} \{Y_i > \sigma \frac{n_u}{n_s}\}\right) \\ &\leq \sum_{i=1}^{n_s} \text{Prob}\left(Y_i > \sigma \frac{n_u}{n_s}\right) \\ &\leq n_s \times n_s^{-\frac{n_s}{3\sigma^2 \log^2(n_s)}} \\ &= n_s^{-\left(\frac{n_s}{3\sigma^2 \log^2(n_s)} - 1\right)}, \end{aligned}$$

which is in the order of  $\frac{1}{n_s^{\Omega(1)}}$ .  $\square$

### 4 Proof for Theorem 6

The proof is similar to that of Theorem 3 in that we create a witness tree, prune it, and then show that a pruned witness tree is unlikely. However, Algorithm MaxBitRate

differs with Algorithm GoWithTheWinner differs in that it's a synchronous algorithm that all users make requests in synchronization and the algorithm executes in discrete time steps rather than continuous time scale. Thus before the proof, we need the following lemmas to assist the formal proof.

LEMMA 7. For any time  $t > 0$ , if user  $u$  receives a application miss from server  $s$  at some time  $t$ , then server  $s$  is overbooked at time  $t - 1$ .

PROOF. If user  $u$  requested a service  $c_u$  from server  $s$  at time  $t$ , it must have also requested  $c_u$  from server  $s$  at time  $t - 1$ . As soon as the request for  $c_u$  was processed at time  $t - 1$ , it was placed server  $s$ . There must have been  $\kappa$  other requests for distinct services that caused the service replacement policy to evict  $c_u$ , resulting in the application miss at time  $t$ . Thus, at least  $\kappa + 1$  distinct services were requested from server  $s$  at time  $t - 1$ , i.e., server  $s$  is overbooked at time  $t - 1$ .  $\square$

To prove convergence, we choose the *sliding window size*  $\tau = c\kappa \log n$ , for a suitably large positive constant  $c$ . Further, consider an initial time interval from time zero to time  $T$  that consists of  $\rho$  intervals of length  $\tau + 1$  each, where  $\rho = 2 \log \log n / \log(\kappa + 1)$ . Thus,  $T = \rho \cdot (\tau + 1) = O(\kappa \log n \log \log n / \log(\kappa + 1))$ .

LEMMA 8. The probability that some user  $u \in U$  decides on an overbooked server  $s \in S$  at some time  $t$ ,  $0 \leq t \leq T$ , is at most  $1/n^{\Omega(1)}$ .

PROOF. Suppose user  $u$  decides on an overbooked server  $s$  at time  $t$ . Then, it must be the case that  $H_\tau(u, s, t) = 100\%$ . Thus, server  $s$  provided a application hit to user  $u$  in every time  $t'$ ,  $t - \tau < t' \leq t$ . Recall that each server serves simultaneous requests by first batching the requests according to the requested applications, i.e., each batch contains requests for the same application, and then serving each batch in random order. Since server  $s$  is overbooked at time  $t$ , it must have been overbooked during all the previous time steps. An overbooked server  $s$  has at least  $\kappa + 1$  distinct applications being requested, i.e., it has at least  $\kappa + 1$  batches of requests. The request made by user  $u$  will receive a application miss if the batch in which it belongs to is  $\kappa + 1$  or higher in the random ordering. Thus, the probability that user  $u$  receives a application miss from the overbooked server  $s$  in any time step  $t' \leq t$  is at least  $1/(\kappa + 1)$ . Since  $H_\tau(u, s, t)$  is 100% only if there is no application miss at any time  $t'$ ,  $t - \tau < t' \leq t$ , the probability of such an occurrence is at most

$$\left(1 - \frac{1}{\kappa + 1}\right)^\tau = \left(1 - \frac{1}{\kappa + 1}\right)^{c\kappa \log n} = \frac{1}{n^{\Omega(1)}},$$

since  $\tau = c\kappa \log n$ . Using the union bound and choosing a suitably large constant  $c$ , the probability that there exists a user  $u \in U$  who decides on an overbooked server  $s$  at some  $t$ ,  $0 \leq t \leq T$ , is at most

$$n \times n \times (T + 1) \times \frac{1}{n^{\Omega(1)}} \leq \frac{1}{n^{\Omega(1)}},$$

since there are  $n$  users, at most  $n$  overbooked servers, and  $T + 1 = O(\kappa \log n \log \log n / \log(\kappa + 1))$  time steps.  $\square$

Now with the two lemmas above, we can prove Theorem 6 as the following.

PROOF. For simplicity, we prove the situation where  $\sigma = 2$ , i.e., each user initially chooses two random candidate servers in step 1 of the algorithm. The case where  $\sigma > 2$  is analogous. Wlog, we also assume  $\kappa$  is at most  $O(\log n / \log \log n)$ , which includes the practically interesting case of  $\kappa$  equal to a constant. When the server capacity is larger, i.e., if  $\kappa = \Omega(\log n / \log \log n)$ , there will be no overbooked servers with high probability and the theorem holds trivially. This observation follows from a well-known result that if  $n$  balls (i.e., users) uniformly and randomly select  $k$  out of  $n$  bins (i.e. servers), then the maximum number of users that select a server is  $O(\log n / \log \log n)$  with high probability, when  $k$  is a fixed constant [19].

In contradiction to the theorem, suppose some user  $u$  (say) has not decided on a server by time  $T$ . We construct a “witness tree<sup>4</sup>” of degree  $\kappa + 1$  and depth at least  $\rho$ , where  $\rho = T/(\tau + 1) = 2 \log \log n / \log(\kappa + 1)$ . Each node of the witness tree is a server. Each edge of the witness tree is a user whose two nodes correspond to the two servers chosen by that user. We show that the existence of an undecided user in time step  $T$  is unlikely by enumerating all possible witness trees and showing that the occurrence of any such witness tree is unlikely. The proof proceeds in the following three steps.

**(1) Constructing a witness tree.** If algorithm MaxHitRate has not converged to the optimal state at time  $T$ , then there exists a user (say  $u_1$ ) and a server  $s$  such that  $H_\tau(u_1, s, T - 1) < 100\%$ , since user  $u_1$  has not yet found a server with a 100% hit rate. We make server  $s$  the root of the witness tree.

We find children for the root  $s$  to extend the witness tree as follows. Since  $H_\tau(u_1, s, T - 1) < 100\%$ , there exists a time  $t'$ ,  $T - 1 - \tau < t' \leq T - 1$ , such that user  $u_1$  received a application miss from server  $s$ . By Lemma 7, server  $s$  was overbooked at time  $t' - 1$ , i.e., there are at least  $\kappa + 1$  users requesting server  $s$  for  $\kappa + 1$  distinct applications at time  $t' - 1$ . Let  $u_1, \dots, u_{\kappa+1}$  be the users who sent requests to server  $s$  at time  $t' - 1$ . Wlog, assume that the users  $\{u_i\}$  are ordered in ascending order of their IDs. By Lemma 8, we know that the probability of a user deciding on an overbooked server is small, i.e., at most  $1/n^{\Omega(1)}$ . Thus, with high probability, users  $u_1, \dots, u_{\kappa+1}$  are undecided at time  $t' - 1$  since they made a request to an overbooked server  $s$ . Let  $s_i$  be the other server choice associated with user  $u_i$  (one of the choices is server  $s$ ). We extend the witness tree by creating  $\kappa + 1$  children for the root  $s$ , one corresponding to each server  $s_i$ . Note that for each of the servers  $s_i$  we know that  $H(u_i, s_i, t' - 2) < 100\%$ , since otherwise user  $u_i$  would have decided on server  $s_i$  in time step  $t' - 2$ . Thus, analogous to how we found children for  $s$ , we can recursively find  $\kappa + 1$  children for each of the servers  $s_i$  and grow the witness tree to an additional level.

Observe that to add an additional level of the witness tree we went from server  $s$  at time  $T - 1$  to servers  $s_i$  at time  $t' - 2$ , i.e., we went back in time by an amount of  $T - 1 - (t' - 2) \leq \tau + 1$ . If we continue the same process, we can construct a witness tree that is a  $(\kappa + 1)$ -ary tree of depth  $T/(\tau + 1) = \rho$ .

**(2) Pruning the witness tree.** If the nodes of the witness tree are guaranteed to represent distinct servers, proving our probabilistic bound is relatively easy. The reason is that if

the servers are unique then the users that represent edges of the tree are unique as well. Therefore the probabilistic choices that each user makes is independent, making it easy to evaluate the probability of occurrence of the tree. However, it may not be the case that the servers in the witness tree constructed above are unique, leading to dependent choices that are hard to resolve. Thus, we create a *pruned witness tree* by removing repeated servers from the original (unpruned) witness tree.

We prune the witness tree by visiting the nodes of the witness tree iteratively in breadth-first search order starting at the root. As we perform breadth-first search (BFS), we remove (i.e., prune) some nodes of the tree and the subtrees rooted at these nodes. What is left after this process is the pruned witness tree. We start by visiting the root. In each iteration, we visit the next node  $v$  in BFS order that has not been pruned. Let  $\beta(v)$  denote the nodes visited *before*  $v$ . If  $v$  represents a server that is different from the servers represented by nodes in  $\beta(v)$ , we do nothing. Otherwise, prune all nodes in the subtree rooted at  $v$ . Then, mark the edge from  $v$  to its parent as a *pruning edge*. (Note that the pruning edges are not part of the pruned witness tree.) The procedure continues until either no more nodes remain to be visited or there are  $\kappa + 1$  pruning edges. In the latter case, we apply a final pruning by removing all nodes that are yet to be visited, though this step does not produce any more pruning edges. This process results in a pruned witness and a set of  $p$  (say) pruning edges.

Note that each pruning edge corresponds to a user who we will call a *pruned user*. We now make a pass through the pruning edges to select a set  $P$  of unique pruned users. Initially,  $P$  is set to  $\emptyset$ . We visit the pruning edges in BFS order and for each pruning edge  $(u, v)$  we add the user corresponding to  $(u, v)$  to  $P$ , if this user is distinct from all users currently in  $P$  and if  $|P| < \lceil p/2 \rceil$ , where  $p$  is the total number of pruning edges. We stop adding pruned users to set  $P$  when we have exactly  $\lceil p/2 \rceil$  users. Note that since a user who made server choices of  $u$  and  $v$  can appear at most twice as a pruned edge, once with  $u$  in the pruned witness tree and once with  $v$  in the pruned witness tree. Thus, we are guaranteed to find  $\lceil p/2 \rceil$  distinct pruned users.

After the pruning process, we are left with a pruned witness tree with nodes representing distinct servers and edges representing distinct users. In addition, we have a set  $P$  of  $\lceil p/2 \rceil$  distinct pruned users, where  $p$  is the number of pruning edges.

**(3) Bounding the probability of pruned witness trees.** We enumerate possible witness trees and bound their probability using the union bound. Observe that since the (unpruned) witness tree is a  $(\kappa + 1)$ -ary tree of depth  $\rho$ , the number of nodes in the witness tree is

$$m = \sum_{0 \leq i \leq \rho} (\kappa + 1)^i = \frac{(\kappa + 1)^{\rho+1} - 1}{\kappa} \leq 2 \log^2 n, \quad (8)$$

since  $\rho = 2 \log \log n / \log(\kappa + 1)$  and hence  $(\kappa + 1)^\rho = \log^2 n$ .

*Ways of choosing the shape of the pruned witness tree.* The shape of the pruned witness tree is determined by choosing the  $p$  pruning edges of the tree. The number of ways of selecting the  $p$  pruning edges is at most  $\binom{m}{p} \leq m^p$ , since there are at most  $m$  edges in the (unpruned) witness tree.

*Ways of choosing users and servers for the nodes and edges of the pruned witness tree.* The enumeration proceeds

<sup>4</sup>A witness tree is so called as it bears witness to the occurrence of an event such as a user being undecided.

by considering the nodes in BFS order. The number of ways of choosing the server associated with the root is  $n$ . Consider the  $i^{\text{th}}$  internal node  $v_i$  of the pruned witness tree whose server has already been chosen to be  $s_i$ . Let  $v_i$  have  $\delta_i$  children. There are at most  $\binom{n}{\delta_i}$  ways of choosing distinct servers for each of the  $\delta_i$  children of  $v_i$ . Also, since there are at most  $n$  users in the system at any point in time, the number of ways to choose distinct users for the  $\delta_i$  edges incident on  $v_i$  is also at most  $\binom{n}{\delta_i}$ . There are  $\delta_i!$  ways of pairing the users and the servers. Further, the probability that a chosen user chooses server  $s_i$  corresponding to node  $v_i$  and a specific one of  $\delta_i$  servers chosen above for  $v_i$ 's children is

$$\frac{1}{\binom{n}{2}} = \frac{2}{n(n-1)},$$

since each set of two servers is equally likely to be chosen in step 1 of the algorithm. Further, note that each of the  $\delta_i$  users chose  $\delta_i$  distinct applications and let the probability of occurrence of this event be  $Uniq(n_c, \delta_i)$ . This uniqueness probability has been studied in the context of collision-resistant hashing and it is known [3] that  $Uniq(n_c, \delta_i)$  is largest when the content popularity distribution is the uniform distribution ( $\alpha = 0$ ) and progressively becomes smaller as  $\alpha$  increases. In particular,  $Uniq(n_c, \delta_i) \leq e^{-\Theta(\delta_i^2/n_c)} < 1$ . Putting it together, the number of ways of choosing a distinct server for each of the  $\delta_i$  children of  $v_i$ , choosing a distinct user for each of the  $\delta_i$  edges incident on  $v_i$ , choosing a distinct application for each user, and multiplying by the appropriate probability is at most

$$\binom{n}{\delta_i} \cdot \binom{n}{\delta_i} \cdot \delta_i! \cdot \left(\frac{2}{n(n-1)}\right)^{\delta_i} \cdot Uniq(n_c, \delta_i) \leq \frac{2^{\delta_i}}{\delta_i!}, \quad (9)$$

provided  $\delta_i > 1$ . Let  $m'$  be the number of internal nodes  $v_i$  in the pruned witness tree such that  $\delta_i = \kappa + 1$ . Using the bound in Equation 9 for only these  $m'$  nodes, the number of ways of choosing the users and servers for the nodes and edges respectively of the pruned witness tree weighted by the probability that these choices occurred is at most

$$n \cdot (2^{\kappa+1}/(\kappa+1)!)^{m'}.$$

*Ways of choosing the pruned users in  $P$ .* Recall that there are  $\lceil p/2 \rceil$  distinct pruned users in  $P$ . The number of ways of choosing the users in  $P$  is at most  $n^{\lceil p/2 \rceil}$ , since at any time step there are at most  $n$  users in the system to choose from. Note that a pruned user has both of its server choices in the pruned witness tree. Therefore, the probability that a given user is a pruned user is at most  $m^2/n^2$ . Thus the number of choices for the  $\lceil p/2 \rceil$  pruned users in  $P$  weighted by the probability that these pruned users occurred is at most

$$n^{\lceil p/2 \rceil} \cdot (m^2/n^2)^{\lceil p/2 \rceil} \leq (m^2/n)^{\lceil p/2 \rceil}.$$

*Bringing it all together.* The probability that there exists a pruned witness tree with  $p$  pruning edges, and  $m'$  internal nodes with  $(\kappa + 1)$  children each, is at most

$$\begin{aligned} & m^p \cdot n \cdot (2^{\kappa+1}/(\kappa+1)!)^{m'} \cdot (m^2/n)^{\lceil p/2 \rceil} \\ & \leq n \cdot (2^{\kappa+1}/(\kappa+1)!)^{m'} \cdot (m^4/n)^{\lceil p/2 \rceil} \\ & \leq n \cdot (2e/(\kappa+1))^{m'(\kappa+1)} \cdot (m^4/n)^{\lceil p/2 \rceil}, \end{aligned} \quad (10)$$

since  $(\kappa + 1)! \geq ((\kappa + 1)/e)^{\kappa+1}$ . There are two possible cases depending on how the pruning process terminates. If

the number of pruning edges,  $p$ , equals  $\kappa + 1$  then the third term of Equation 10 is

$$(m^4/n)^{\lceil p/2 \rceil} \leq (16 \log^8 n/n)^{\lceil (\kappa+1)/2 \rceil} \leq 1/n^{\Omega(1)},$$

using Equation 8 and assuming that cache size  $\kappa$  is at least a suitably large constant. Alternately, if the pruning process terminates with fewer than  $\kappa + 1$  pruning edges, it must be that at least one of the  $\kappa + 1$  subtrees rooted at the children of the root  $s$  of the (unpruned) witness tree has no pruning edge. Thus, the number of internal nodes  $m'$  of the pruned witness tree with  $(\kappa + 1)$  children each is bounded as follows:

$$m' = \sum_{0 \leq i < \rho-1} (\kappa+1)^i \geq (\kappa+1)^{\rho-2} \geq \log^2 n/(\kappa+1)^2,$$

as  $(\kappa + 1)^\rho = \log^2 n$ . Thus, the second term of Equation 10 is

$$(2e/(\kappa+1))^{m'(\kappa+1)} \leq (2e/(\kappa+1))^{\log^2 n/(\kappa+1)} \leq 1/n^{\Omega(1)},$$

assuming  $\kappa > 2e - 1$  but is at most  $O(\log n / \log \log n)$ . Thus, in either case, the bound in Equation 10 is  $1/n^{\Omega(1)}$ . Further, since there are at most  $m$  values for  $p$ , the total probability of a pruned witness tree is at most  $m \cdot 1/n^{\Omega(1)}$  which is  $1/n^{\Omega(1)}$ . This completes the proof of the theorem.  $\square$